

Sketch-based Video Object Localization

Sangmin Woo¹ So-Yeong Jeon^{1,2} Jinyoung Park¹ Minji Son³ Sumin Lee¹ Changick Kim¹
¹KAIST ²Korea Agency for Defense Development ³LG Electronics

¹{smwoo95, presentover, jinyoungpark, suminlee94, changick}@kaist.ac.kr ³minji13.son@lge.com

<https://github.com/sangminwoo/SVOL>

Abstract

We introduce Sketch-based Video Object Localization (SVOL), a new task aimed at localizing spatio-temporal object boxes in video queried by the input sketch. We first outline the challenges in the SVOL task and build the Sketch-Video Attention Network (SVANet) with the following design principles: (i) to consider temporal information of video and bridge the domain gap between sketch and video; (ii) to accurately identify and localize multiple objects simultaneously; (iii) to handle various styles of sketches; (iv) to be classification-free. In particular, SVANet is equipped with a Cross-modal Transformer that models the interaction between learnable object tokens, query sketch, and video through attention operations, and learns upon a per-frame set matching strategy that enables frame-wise prediction while utilizing global video context. We evaluate SVANet on a newly curated SVOL dataset. By design, SVANet successfully learns the mapping between the query sketches and video objects, achieving state-of-the-art results on the SVOL benchmark. We further confirm the effectiveness of SVANet via extensive ablation studies and visualizations. Lastly, we demonstrate its transfer capability on unseen datasets and novel categories, suggesting its high scalability in real-world applications.

1. Introduction

A sketch is worth a thousand words. It can even convey ideas that are hard to explain in words. Due to the concise and abstract nature of the sketch, it can be illustrative, making it an excellent tool for a variety of applications [3, 10, 32, 38, 57, 81, 86, 90]. Meanwhile, query-based localization is one of the long-sought goals for visual understanding. The literature has been studied at a variety of query types (e.g., image, language, sketch) and domains (e.g., image, video) [1, 11, 30, 36, 42, 49, 65, 71, 74, 83, 89]. While numerous studies have shown remarkable results using image or language as query, both have their own limitations. Images containing a specific object of interest may be difficult to collect due to privacy or copyright issues [3], and

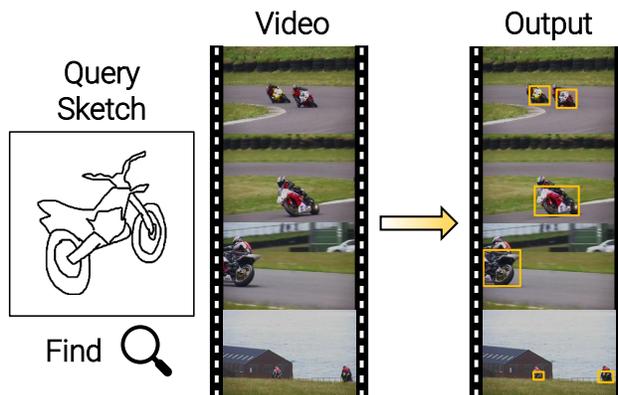


Figure 1. **Illustration of the SVOL task.** Given a query sketch, the goal is to find all *object boxes* (colored in yellow) spatio-temporally that match the sketch object in a video. Query sample is randomly drawn from *Sketchy* dataset.

the utility of language is limited as it varies per country. As an alternative, using sketch as a query brings several advantages. It allows for immense expressive flexibility and can transcend language barriers [47]. Moreover, due to the recent spread of touchscreen devices (e.g., smartphones, tablets), sketches have become easier to obtain than ever [73]. Amid the explosive growth of video data, sketch has emerged as an appealing candidate for user interface in online video platforms thanks to these properties. Despite its promise, the use of sketch as query for object localization in the video domain has not yet been explored.

In this work, we propose a new task called Sketch-based Video Object Localization (SVOL) that aims to localize objects in videos with the query sketch (see Fig. 1). We first identify several challenges that the setting of SVOL yields, including but not limited to: (i) As objects move, they can generate motion blurs or occlude parts of other objects, thus distorting their appearances [45]. Moreover, objects in the scene may suddenly disappear, or objects that were not in the scene may suddenly appear. These dynamic changes over time complicate the matching of sketch to its corresponding objects. (ii) Multiple objects can appear in a video. Therefore, it is important not only to accurately differentiate between

the target objects from multiple objects belonging to different categories, but also to find all objects that match the sketch query simultaneously. **(iii)** As for sketch, a single object can be drawn in various ways [67]. Unlike natural videos, sketches lack color, texture, and background information, resulting in a high degree of freedom. This allows sketches to be drawn in a variety of styles (*i.e.*, different abstraction levels). **(iv)** There should be no explicit category prediction (*i.e.*, no fixed classes) in the SVOL system. As with all query-based localization tasks, SVOL requires finding the best matching objects given a query sketch (not the category itself).

Driven by this analysis, we propose SVANet that serves as a strong baseline for the SVOL task. SVANet takes extracted video and sketch representations as inputs and predicts box coordinates and objectness scores end-to-end. Our SVANet is built on several design principles: **First**, we propose a novel *Cross-modal Transformer (CMT)* that not only closes the domain discrepancy between sketch and video but also models video temporal context. We equip CMT with four attention operations [76] to leverage their strong relational modeling capability. By design, CMT emphasizes important content by learning the correlation between sketch and video representations, and incorporates temporal context by modeling intra-content relationships. Also, CMT takes object tokens as inputs and transforms them into predictions of box coordinates and objectness scores by learning their internal interactions and by referring to joint sketch-video representations. **Second**, we formulate the SVOL task as a set prediction problem [8] and employ a *per-frame set matching* strategy. We predict all bounding boxes across the video frames, and find the best matching between predicted and ground truth boxes that minimizes the matching cost. The overall training loss is then defined based on the matching results. Instead of matching whole video-level results with video-level ground truths, we perform set matching frame-by-frame. This enables the prediction of multiple objects in parallel while utilizing the global video context. **Third**, SVANet is designed to be compatible with various sketch styles. SVANet learns to embed the sketch objects of the same category into a similar subspace of a high-dimensional latent space, regardless of differences in sketch styles (*e.g.*, shape, pose, line thickness, *etc.*). This style-agnostic property enables SVANet to generalize well on unseen sketch datasets with varying degrees of abstraction. **Last**, SVANet has no explicit classification in the pipeline. This allows SVANet to learn the mapping between sketch and video objects based on implicit similarity (*e.g.*, symbolic meaning, appearance, *etc.*). This classification-free property of SVANet extends its applicability to any kind of free-form sketches, allowing us to query over novel object classes.

To benchmark our approach and show the potential of using sketch as query, we present a new SVOL dataset curated

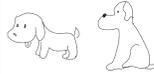
	Sketchy	TU-Berlin	QuickDraw
cat			
dog			
Abstract	Low	Medium	High

Figure 2. **Sketch datasets comparison.** **Sketchy** [69] is the most realistic since it is drawn after photographic objects, **QuickDraw** [34] has the highest level of abstraction due to limited drawing time (< 20 secs), and **TU-Berlin** [20] lies halfway between them.

from the video dataset, ImageNet-VID [66], and three sketch datasets with varying degrees of abstraction (see Fig. 2): **Sketchy** [69], **TU-Berlin** [20], and **QuickDraw** [34]. We show that SVANet outperforms the strong image-level baseline (Sketch-DETR) [65] by a significant margin: $\sim 29.4\%$, 17.7% , and 16.8% improvement of mIoU using **Sketchy**, **TU-Berlin**, and **QuickDraw** sketch datasets, respectively. This implies that SVANet effectively resolves the limitation of image-level baselines with temporal video context. Moreover, we verify the effectiveness of several design choices of SVANet through extensive ablation studies and analyze its behavior with several visualizations. Finally, we evaluate transfer performances of SVANet on unseen datasets (with different abstraction levels or sketch styles) and novel categories that are unseen during training. The results demonstrate that SVANet is robust to style variations and that the learned sketch-video mapping function generalizes well to novel classes of sketches. These appealing properties are ideal for several query-based applications in practice, such as large-scale video platforms, in that the system can flexibly respond to diverse inputs from users.

Our contributions are summarized as follows:

New task. We opened up a new challenging task, Sketch-based Video Object Localization (SVOL). We also identified several challenges that the SVOL task setting brings.

New dataset. We presented a new SVOL dataset curated from the video dataset (ImageNet-VID [66]) and three sketch datasets with different styles (**Sketchy** [69], **TU-Berlin** [20], **QuickDraw** [34]) and provide a benchmark with comparison against the frame-level baselines and several variants.

Strong baseline. We proposed a novel framework named SVANet, equipped with SVOL-tailored designs such as Cross-modal Transformer and per-frame set matching, that serves as a strong baseline on the benchmark: SVANet improves mIoU by 29.4% , 17.7% , 16.8% over the strong counterpart, Sketch-DETR [65], using **Sketchy**, **TU-Berlin**, **QuickDraw** sketch dataset, respectively.

Extensive experiments. We thoroughly investigated the effects of model components with comprehensive ablative studies in various aspects. Last but not least, we demonstrated the strong generalizability of SVANet on unseen datasets

and novel categories, which makes sketch as query highly practical in real-world scenarios.

2. Related Work

Our work builds on previous work in several areas, including sketch-based applications, query-based localization, and Transformer architecture.

2.1. Sketch-based Applications

Our work builds on the idea of using sketches as a way to query visual data. Sketch is a universal communication tool that is not bound by age, race, language, or national boundaries. Recently, sketch-based applications have grown at an unprecedented rate due to the widespread use of touch-screen devices such as smartphones and tablets that enable acquiring sketch data much easier than ever. Here are some examples of various applications using sketch ¹:

Image retrieval: a user sketches an object or scene and the system retrieves similar image from a database [3, 18, 67, 90].

Image synthesis: a user sketches an image and the system generates a photorealistic version of the image [10, 32, 81];

Image editing: a user sketches desired changes to an image, and the system makes the changes automatically [57, 88].

Robot interface: a user sketches a task for a robot to perform, such as picking up an object and placing it in a specific location [5, 68].

3D modeling: a user sketches a 3D object and the system generates a 3D model of the object [54, 78].

Augmented reality: a user sketches an object, and the system overlays a 3D model of the object in the real-world environment [33, 38].

Additionally, sketch is particularly effective at representing detailed features of an object like its shape, pattern, and pose. This ability to convey such fine-grained information has made sketches a popular tool in a variety of studies, including image [91], scene [48], and video [87] retrieval. For example, in fine-grained image retrieval, sketch are used as queries to retrieve specific objects within images, such as a specific breed of dog or type of car [70, 72]. This allows users to search for images with specific visual characteristics, such as the shape of a dog’s ears or a car’s grille, so that objects of the same category can be differentiated.

In the SVOL problem, while sketches have the capability to provide fine-grained information, we opt to focus on category-level object localization, *i.e.*, localization is carried out in a shape- and pose-agnostic manner within the same category. This is because it is not natural to match a static sketch that has a specific shape and pose with objects in a video, whose shape and pose dynamically change over time. Additionally, by focusing on category-level localization, we

¹For a more detailed list of sketch-based applications, we recommend referring to [86].

can take advantage of the abstract nature of sketches. We can identify the location of objects in a video by sketching only some key features of the object, such as the headlights and grille of a car.

2.2. Query-based Localization

SVOL is related to the literature on object detection and tracking in videos, with added constraint of using a sketch as the query. Query-based object localization is similar to object detection [46, 50, 62, 63] (or video object detection [12, 23, 28, 79, 84, 94]) in that they both aim to locate the bounding boxes of objects in an image (or a video). However, query-based localization is grounded on the given query, rather than pre-defined object classes. Query-based localization tasks have been studied using various query types in diverse dimensions.

Query. Image queries can be localized based on appearance similarity, allowing themselves to be easily transferred to other objects with just a few image samples. This desirable property opens up a new avenue for research on one/few-shot localization [22, 30, 55, 77]. However, image queries are hard to acquire in some privacy or security-related situations, making their usage in some applications difficult.

Language queries, on the other hand, are highly useful given that we just need to describe the objects of interest in natural language. However, its universality is limited since the assumed language (English) may not be familiar to some people (non-native English speakers). As such, when the language is re-targeted, the neural network may require extra learning or translation before providing the query. **Sketch** queries differ from image queries in that they lack rich information such as color, texture, and background information; most free-hand sketches are composed solely of monochromatic lines, with no texture and context. In addition, since the sketches are drawn by envisioning abstract objects, even the same object may be drawn differently by a different person. These characteristics make sketch-based localization challenging. Nevertheless, we argue that this line of research is valuable since it offers the highest degree of freedom among the three query types and can transcend the language barrier (*i.e.*, the sketch of ‘cat’ can be understood whether or not you are a native English speaker). Our work focuses on the emerging role of sketch in the context of query-based localization.

Dimension. Temporal localization aims to identify the temporal span (1D) in which the query object appears in the video. **Spatial** localization attempts to locate all object instances that match the query object within a still image (2D). **Spatio-temporal** localization seeks to locate all object instances that match the query object in every frame of video (3D). Our work belongs to the spatio-temporal category.

Task. The challenging and open-ended nature of the query-based localization problem lends itself to a variety of tasks:

image-based localization in natural images [2, 14, 49, 52] and videos [11, 13, 41, 42] (a.k.a., visual object tracking); language-based localization in natural images [15, 16, 24, 31, 36, 89] (a.k.a., visual grounding or referring expression comprehension) and videos [1, 6, 21, 27, 40, 71, 83, 93] (a.k.a., video grounding or natural language video localization); and sketch object localization in natural images [65, 74]. The most relevant tasks to ours are video grounding [21, 40, 71, 93] and sketch object localization in images [65, 74].

Comparison. In the realm of query-based localization research, various query types and domains have been explored, such as image, language, and video. However, one noticeable gap in the existing literature pertains to the absence of studies focused on sketch queries in the video domain. We seek to address this particular gap in knowledge. To address this research void, we propose a novel task “Sketch-based Video Object Localization”. This task is designed to facilitate the precise localization of spatio-temporal object boxes within video content, with the query input provided in the form of a sketch. This novel approach bridges the gap between sketch-based queries and video object localization, opening up new avenues for exploration and advancement in the field.

2.3. Sketch-based Image Object Localization

There are few methods for image object localization using sketch queries [65, 74], and we are the first to propose a sketch-based video object localization approach. We adopt Cross-modal Attention [74] and Sketch-DETR [65] as the image-level baseline in our SVOL benchmark.

Tripathi *et al.*, Cross-modal Attention [74] generates object proposals that match the query sketch in an image. This mechanism operates above the off-the-shelf object detection framework, Faster R-CNN [63], in which the key component is region proposal network (RPN). Tripathi *et al.* modifies the RPN structure to integrate the sketch information in order to create object proposals that are relevant to the query sketch. In more detail, feature vectors of different regions in the image feature map are scored using the global sketch representation to determine compatibility. The attention feature is then calculated by multiplying these compatibility scores with image feature maps. These attention feature maps are concatenated with the original feature maps and projected to a lower-dimensional space, which is then input to RPN to yield relevant object proposals. The pooled object proposals are scored using a sketch feature vector to localize the object of interest.

Riba *et al.*, Sketch-DETR [65] is built on the DETR [8] architecture. Given a natural image and a query sketch, Sketch-DETR [65] transforms each input with a separate CNN backbone, and generates feature maps for each input modality. They are then fused via concatenation. Specifically, the sketch is inflated by the resolution of the image

feature map then projected using a 1×1 convolution operation. The obtained feature map are flattened before being fed into the Transformer encoder-decoder. The final bounding boxes and their respective score are predicted through a shared feed-forward network (FFN).

2.4. Vision and Multimodal Transformers

Transformer [76] is a universal sequence processor with an attention-based architecture that is originally designed for machine translation. The primary components of Transformer are self-attention that captures long-range interactions within a single context and cross-attention that considers token correspondences between two sequences.

Vision Transformers. Beyond natural language processing [7, 17, 59, 60], Transformers have rapidly become the *de facto* standard in a variety of computer vision applications: image recognition [19], object detection [8], panoptic segmentation [80], human object interaction [35], action recognition [39, 82], and object tracking [11]. Among these, it is worth noting that Detection Transformer (DETR) [8] has made a significant breakthrough in the field of object detection by successfully adopting Transformer design and bipartite matching algorithm. By design, DETR eliminates the need for heuristics (*e.g.*, non-maximum suppression) in the detection pipeline while leveraging the capability of global relation modeling. Inspired by the recent successes of Transformers, particularly DETR, we view the SVOL task as a set prediction problem and build our model on the Transformer architecture. Furthermore, since videos are a sequence of frames, the Transformer is well-suited to model temporal information of videos.

Multimodal Transformers. Transformers have shown to be particularly effective in multimodal processing due to their ability to selectively attend to relevant information from multiple modalities (*e.g.*, text, image, and audio). This has been demonstrated in various tasks, including image captioning [58], visual question answering [9], natural language video grounding [83], text-to-image synthesis [61], and text-to-speech [43]. Recent studies such as CLIP [58] and DALL-E [61] highlighted the potential of pre-training Transformer-based models on a vast amount of image-text pairs using a contrastive loss. This provides a strong starting point when fine-tuning on downstream tasks, thereby allowing the model to generalize well unseen datasets. Furthermore, Transformers have shown to transfer well across tasks, making them a versatile model for various multimodal processing tasks [44, 53]. These properties of Transformer make itself a strong candidate for the SVOL task, where the intricate relationship between the query sketch and video objects must be modeled to bridge the gap between natural video and sketches with various styles.

3. Preliminary: Transformer

We build our Cross-modal Transformer (CMT) on top of Transformer design², which: (i) densely relates every pair of elements in the input sequence; (ii) captures long-range context with minimal inductive bias (compared to CNNs or RNNs); (iii) effectively models interaction between multi-modal (cross-domain) sequences. These desirable properties of Transformer makes itself well-suited for the CMT design.

The common practice is to use attention mechanism with residual connection, dropout, and layer normalization. Attention in the general QKV form is a popular yet strong mechanism for neural systems. Given that attention operations are key building blocks of CMT, we first briefly discuss their general form.

3.1. QKV Attention

Given input sequences $\mathbf{q} \in \mathbb{R}^{L_1 \times D}$, $\mathbf{k} \in \mathbb{R}^{L_2 \times D}$ and $\mathbf{v} \in \mathbb{R}^{L_2 \times D}$, we project them into separate embedding spaces. We call the embedded representations as query (\mathbf{Q}), key (\mathbf{K}), and value (\mathbf{V}).

$$\mathbf{Q} = (\mathbf{q} + \text{pos}_q) \mathbf{W}_q, \quad (1)$$

$$\mathbf{K} = (\mathbf{k} + \text{pos}_k) \mathbf{W}_k, \quad (2)$$

$$\mathbf{V} = \mathbf{v} \mathbf{W}_v, \quad (3)$$

where $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v \in \mathbb{R}^{D \times D_h}$ are learnable weights. Since the Transformer is inherently permutation-invariant w.r.t input sequence, we add positional encoding $\text{pos}_q \in \mathbb{R}^{L_1 \times D}$ and $\text{pos}_k \in \mathbb{R}^{L_2 \times D}$ (fixed absolute encoding to represent positions using sine and cosine functions of different frequencies) to embedded sequences.

The attention weights \mathbf{A}_{ij} are computed by comparing two elements of the sequence (dot products) to their respective query \mathbf{Q}_i and key \mathbf{K}_j representations, normalized by D_h .

$$\mathbf{A} = \text{softmax} \left(\frac{\mathbf{QK}^T}{\sqrt{D_h}} \right), \quad (4)$$

Finally, we calculate a weighted sum over all value representation \mathbf{V} .

$$\text{Att}(\mathbf{q}, \mathbf{k}, \mathbf{v}) = \mathbf{AV}. \quad (5)$$

We call this operation as *Self-Attention* for the special case where \mathbf{q} , \mathbf{k} , and \mathbf{v} are all the same.

3.2. Multi-Head Attention

Multi-Head Attention (MHA) allows the model to jointly attend to information from different representation subspaces at different positions. It is a simple extension of Attention in which several Attention heads are executed in parallel followed by a projection of their concatenated outputs. To

maintain the computed value and the number of parameters constant when changing the number of heads k , D_h is typically set to D/k .

$$\text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v}) = [\text{Att}_1(\mathbf{q}, \mathbf{k}, \mathbf{v}); \dots; \text{Att}_k(\mathbf{q}, \mathbf{k}, \mathbf{v})] \mathbf{W}_{\text{MHA}}, \quad (6)$$

where $[\cdot]$ denotes concatenation on the channel axis and $\mathbf{W}_{\text{MHA}} \in \mathbb{R}^{k \cdot D_h \times D}$ is learnable weight.

4. Method

We begin by describing the SVOL task and present an end-to-end trainable SVANet that predicts a set of objects based on dense pair-wise relation modeling. Next, we introduce a per-frame set matching strategy that imposes a unique match between predicted and ground truth boxes at each frame; then, define an overall training loss. An overview of SVANet is depicted in Fig. 3.

SVOL task definition. Given a query sketch \mathcal{S} and a video \mathcal{V} , the goal of SVOL is to find all spatio-temporal boxes \mathcal{Y} that *match* the sketch object in the video. We consider the video as a sequence of L frames, $\mathcal{V} = [V_i]_{i=1}^L$, and aim to find all boxes $\mathcal{Y} = [B_i]_{i=1}^L$ over the video frames, where $B_i \in \mathbb{R}^{K_i \times 4}$ is a set of bounding boxes at video frame V_i , $B_i = \{b_i^j\}_{j=1}^{K_i}$. The number of boxes K_i at frame V_i can vary throughout the video, since objects can be occluded, disappear or appear in the scene. We predict a total of N bounding boxes across L frames, $\hat{\mathcal{B}} = [\hat{B}_i]_{i=1}^L$, M boxes per frame, $\hat{B}_i = \{\hat{b}_i^j\}_{j=1}^M$, where $N = L \times M$. The predictions are considered as *correct* if IoU between the predicted box \hat{b} and the ground truth box b is higher than the threshold μ . The bounding box is defined as a 4D vector normalized w.r.t. the frame resolution: $b \in [0, 1]^4$. We also predict the likelihoods that the predicted boxes contain the target object, referred to as objectness scores $\hat{\mathcal{O}}$, where each element $\hat{o} \in [0, 1]$. In short, the predictions are a set of bounding boxes and their corresponding objectness scores: $\hat{\mathcal{Y}} = \{\hat{\mathcal{B}}, \hat{\mathcal{O}}\}$. As we view SVOL as a set prediction problem, we find the best matching between the ground truth set \mathcal{Y} and the prediction set $\hat{\mathcal{Y}}$.

As we set the SVOL problem as category-level localization, the system is trained to perform the consistent bounding box localization for sketches belonging to the same category, regardless of variations in shape or pose. This allows the system to operate robustly, even in the presence of different levels of abstraction or diverse styles in the sketches. However, it is worth noting that there is *no* explicit category prediction inside the system, instead it relies on implicit similarity (*e.g.*, symbolic meaning, appearance, *etc.*) to learn sketch-video object matching.

4.1. SVANet Architecture

SVANet is designed to address the challenge of bridging the gap between two distinct modalities, sketch and natural

²We leave an original paper as a reference [76] for further details of Transformer building blocks.

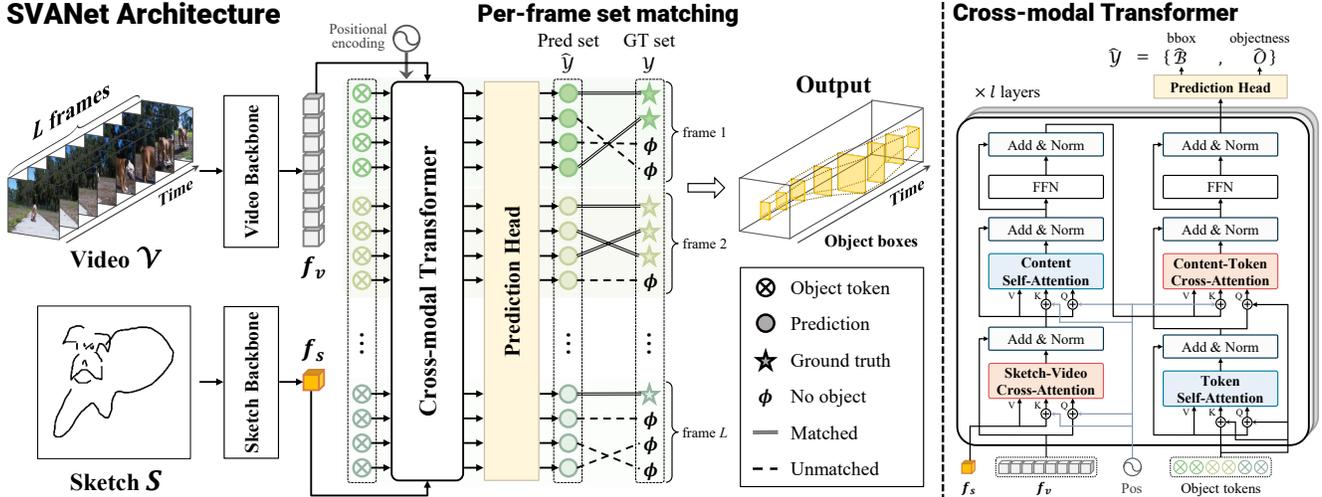


Figure 3. **Overview of SVANet.** Given a video \mathcal{V} and a query sketch \mathcal{S} , SVANet processes them in a separate encoding pipeline, yielding a sequence of frame representations f_v and a sketch representation f_s . The Cross-modal Transformer (CMT) then takes f_v , f_s , and a set of learnable object tokens as inputs. Through the CMT layers, object tokens learn interactions between themselves and attend to sketch-video joint representation to produce accurate predictions. **Per-frame set matching.** During training, SVANet finds the best matching that minimizes the matching cost (see Eq. (15)) between the prediction set and the ground truth set for each frame. To assign a unique matching between the two sets, the ground truth set is padded with additional No object (ϕ) elements. The overall loss is defined based on the matching results (see Eq. (17)). As a result, SVANet outputs the spatio-temporal object boxes. **Cross-modal Transformer.** In CMT, sequences of representations are added with positional encoding before every attention operation. CMT first highlights important contents by learning the correspondence between sketch f_s and video f_v representations, and models intra-content relationships. CMT then transforms the object token set into a set of predictions (box coordinates and objectness scores) by learning token-token interactions and referring to joint sketch-video representations. More details are in Sec. 4. Best viewed in color.

video, in order to perform object localization. The system incorporates attention operations that consider a wide range of contexts and inter-dependencies between elements within the input sequences. This leads our system to acquire the capability to learn powerful representations of the input sequences and delivers accurate video object localization using sketches as queries.

Video & sketch backbones. A video, represented by a sequence of frames, $\mathcal{V} \in \mathbb{R}^{L \times C_0 \times H_0 \times W_0}$, where $L = 32$, $C_0 = 3$, $H_0 = W_0 = 224$, is initially processed using the ResNet-50 architecture [29], generating high-dimensional feature maps $f_v \in \mathbb{R}^{L \times C \times H \times W}$, where $C = 512$, $H = W = 7$. Likewise, a sketch \mathcal{S} is processed using ResNet-18 [29], followed by a spatial pooling operation that compresses it into 1D representation $f_s \in \mathbb{R}^C$. Finally, the outputs f_v and f_s are passed through the Cross-modal Transformer. To address the sparse nature of sketch information, we use a lighter CNN backbone (ResNet-18) compared to the video (ResNet-50).

Cross-modal Transformer & prediction head. In addition to f_v and f_s , the Cross-modal Transformer (CMT)³ takes a set of N learnable embeddings initialized with random weights, which we refer to as *object tokens*, and transforms them into a set of N predictions.

³More details about original Transformer are in Sec. 3.

CMT consists of l layers, and each layer contains four attention operations: (i) *Sketch-Video Cross-Attention (SVCA)* assigns higher attention weights to the important elements of the input sequence (video patches), that are relevant for accurate bounding box localization based on the input sketch query. This is achieved by modeling the inter-modality relationship between video f_v and sketch f_s representations. SVCA bridges the gap between sketches and videos by effectively integrating the information from both modalities. (ii) *Content Self-Attention (CSA)* is responsible for modeling the temporal relationship between the elements in the input sequence (i.e., output of SVCA). By considering the pairwise relationship of these elements, CSA enables a more comprehensive understanding of the broader video context. (iii) *Token Self-Attention (TSA)* receives object tokens as input and models interactions between them, enabling them to globally reason about all objects. (iv) *Content-Token Cross-Attention (CTCA)* transforms object tokens to meaningful outputs by relating them with contextual representation of content (i.e., output of CSA). Since the attention operations are permutation-invariant (i.e., produce the same output regardless of the order of elements in the input sequence), we supplement the input sequence with temporal order information by adding absolute positional encoding prior to every attention operation (except for TSA; we instead add object tokens to the input sequence of each TSA operation).

All attention operations in CMT are in the form of MHA

(Eq. (6)) with 8 heads (*i.e.*, $k = 8$). Let a video representation $f_v \in \mathbb{R}^{L \times C \times H \times W}$ as $\mathbf{v}^{(0)}$ and a sketch representation $f_s \in \mathbb{R}^{1 \times C}$ as \mathbf{s} . Given $\mathbf{v}^{(0)}$ and \mathbf{s} the i -th CMT layer calculates:

$$\mathbf{x}^{(i)} = \text{LN}(\text{SVCA}^{(i)}(\mathbf{v}^{(i)}, \mathbf{s}, \mathbf{s}) + \mathbf{v}^{(i)}), \quad (7)$$

$$\mathbf{y}^{(i)} = \text{LN}(\text{CSA}^{(i)}(\mathbf{x}^{(i)}, \mathbf{x}^{(i)}, \mathbf{x}^{(i)}) + \mathbf{x}^{(i)}), \quad (8)$$

$$\mathbf{v}^{(i+1)} = \text{LN}(\text{FFN}_1^{(i)}(\mathbf{y}^{(i)}) + \mathbf{y}^{(i)}), \quad (9)$$

$$\mathbf{p}^{(i)} = \text{LN}(\text{TSA}^{(i)}(\mathbf{r}^{(i)}, \mathbf{r}^{(i)}, \mathbf{r}^{(i)}) + \mathbf{r}^{(i)}), \quad (10)$$

$$\mathbf{q}^{(i)} = \text{LN}(\text{CTCA}^{(i)}(\mathbf{p}^{(i)}, \mathbf{v}^{(i+1)}, \mathbf{v}^{(i+1)}) + \mathbf{p}^{(i)}), \quad (11)$$

$$\mathbf{r}^{(i+1)} = \text{LN}(\text{FFN}_2^{(i)}(\mathbf{q}^{(i)}) + \mathbf{q}^{(i)}), \quad (12)$$

where LN is layer normalization and FFN is 2-layer feed-forward network. Here, $r^{(0)} = O_{N \times C}$ ($N \times C$ -sized zero matrix), thus TSA operation Eq. (10) can be omitted in the first CMT layer.

TSA (Eq. (10)) and CTCA operations (Eq. (11)) slightly differ with the standard QKV attention in that they consider the object tokens `tkn` as learnable positional encoding for the query (\mathbf{q}) inputs, *i.e.*, \mathbf{q} is added with `tkn` instead of `posq` in Eq. (1).

$$\mathbf{Q} = (\mathbf{q} + \text{tkn})\mathbf{W}_{\mathbf{q}}, \quad (13)$$

In addition, since TSA is Self-Attention operation ($\mathbf{q} = \mathbf{k} = \mathbf{v} = \mathbf{r}^{(i)}$), `tkn` is also used as positional encoding for the key (\mathbf{k}) input in TSA.

$$\mathbf{K} = (\mathbf{k} + \text{tkn})\mathbf{W}_{\mathbf{k}}, \quad (14)$$

The subsequent processes are the same as standard QKV attention.

We go through l CMT layers, and the final CMT output $r^{(l)}$ is fed into two separate linear layers (*i.e.*, prediction heads) to obtain a set of bounding box coordinates $\hat{\mathcal{B}}$ and objectness scores $\hat{\mathcal{O}}$, respectively.

4.2. SVOL as a Set Prediction

In this work, we formulate SVOL as a set prediction problem. In practice, we adopt a Hungarian algorithm [37] to find an optimal matching between predictions and ground truths in a way that minimizes the matching cost. The overall loss function is defined based on the matching results.

Per-frame set matching. SVANet transforms N object tokens to N predictions (bounding boxes and objectness scores). Here, we make each of the N/L (hereafter M) tokens to be responsible for predicting the results of each frame V_i by performing per-frame set matching. This allows SVANet to predict results per frame while being able to access global context information across the video. We formally describe the process in the following.

A set of ground truth bounding boxes \mathcal{Y} can be seen as a sequence of L subsets, where i -th subset has K_i elements:

$[\{b_i^j\}_{j=1}^{K_i}]_{i=1}^L$. Likewise, we evenly divide a prediction set $\hat{\mathcal{Y}} = \{\hat{\mathcal{B}}, \hat{\mathcal{O}}\}$ of size N into L subsets having M elements each: $[\{\hat{y}_i^j\}_{j=1}^M]_{i=1}^L$, where $\hat{y}_i^j = (\hat{b}_i^j, \hat{o}_i^j)$. Hereafter, we denote the i -th subset of ground truths as Y_i and that of predictions as \hat{Y}_i for conciseness. The size of \hat{Y}_i is assumed to be larger than the size of Y_i : $M > K_i$. Since the Hungarian algorithm pairs the elements of two sets one by one, we pad Y_i with `No object` (\emptyset) to match the size of M . For every single i (from $i = 1$ to $i = L$), we seek for the best one-to-one matching between Y_i and \hat{Y}_i using a Hungarian algorithm. Formally, in the i -th prediction subset, let $\hat{y}_i^{\sigma_i(j)}$ be the j -th element under a permutation of M elements $\sigma_i \in \mathfrak{S}_i(M)$. We now define the pair-wise matching cost \mathcal{C} as:

$$\mathcal{C}(b_i^j, \hat{y}_i^{\sigma_i(j)}) = -\mathbb{1}_{\{b_i^j \neq \emptyset\}} \hat{o}_i^{\sigma_i(j)} + \mathbb{1}_{\{b_i^j \neq \emptyset\}} \mathcal{L}_{\text{box}}(b_i^j, \hat{b}_i^{\sigma_i(j)}). \quad (15)$$

Note that the `No object` paddings in the ground truth are not considered when calculating the matching cost. For every i , we aim to find the optimal assignment $\sigma_i^* \in \mathfrak{S}_i(M)$ that pairs the predictions and ground truths at the lowest cost:

$$\sigma_i^* = \underset{\sigma_i \in \mathfrak{S}_i(M)}{\text{argmin}} \sum_{j=1}^M \mathcal{C}(b_i^j, \hat{y}_i^{\sigma_i(j)}). \quad (16)$$

Overall loss. Based on the matching results, our set prediction loss $\mathcal{L}_{\text{set}}(\mathcal{Y}, \hat{\mathcal{Y}})$ is defined as:

$$\sum_{i=1}^L \sum_{j=1}^M \left[-\lambda_{\text{obj}} \log \hat{o}_i^{\sigma_i^*(j)} + \mathbb{1}_{\{b_i^j \neq \emptyset\}} \mathcal{L}_{\text{box}}(b_i^j, \hat{b}_i^{\sigma_i^*(j)}) \right], \quad (17)$$

where $\lambda_{\text{obj}} \in \mathbb{R}$ is a loss coefficient for objectness scores. Here, the log-probability of the `No object` paddings (\emptyset) is scaled down by a factor of 10 to strike a balance between object and no-object.

The box loss \mathcal{L}_{box} is defined as a linear combination of ℓ_1 loss and the generalized IoU (gIoU) loss [64]:

$$\mathcal{L}_{\text{box}}(b_i^j, \hat{b}_i^{\sigma_i(j)}) = \lambda_{\ell_1} \mathcal{L}_{\ell_1}(b_i^j, \hat{b}_i^{\sigma_i(j)}) + \lambda_{\text{iou}} \mathcal{L}_{\text{iou}}(b_i^j, \hat{b}_i^{\sigma_i(j)}), \quad (18)$$

where $\lambda_{\ell_1} \in \mathbb{R}$ and $\lambda_{\text{iou}} \in \mathbb{R}$ are balancing hyperparameters. While both losses have the same goal, object localization, the ℓ_1 loss will have different scales for small and large boxes even if their relative errors are similar, whereas the gIoU loss is scale-invariant.

We calculate the ℓ_1 loss as:

$$\mathcal{L}_{\ell_1}(b_i^j, \hat{b}_i^{\sigma_i(j)}) = \|b_i^j - \hat{b}_i^{\sigma_i(j)}\|_1. \quad (19)$$

The gIoU loss is calculated as (we denote the area with set

operations for the sake of argument):

$$\mathcal{L}_{iou}(b_i^j, \hat{b}_i^{\sigma(j)}) = 1 - \left(\frac{|b_i^j| \cap |\hat{b}_i^{\sigma(j)}|}{|b_i^j| \cup |\hat{b}_i^{\sigma(j)}|} - \frac{|B(b_i^j, \hat{b}_i^{\sigma(j)})| \setminus |b_i^j| \cup |\hat{b}_i^{\sigma(j)}|}{|B(b_i^j, \hat{b}_i^{\sigma(j)})|} \right), \quad (20)$$

where $|\cdot|$ represents the bounding box area, and the symbols \cup , \cap , and \setminus calculate the area of union, intersection, and subtraction of the two bounding box areas, respectively. $B(b_i, \hat{b}_{\sigma(i)})$ denotes the smallest box enclosing b_i and $\hat{b}_{\sigma(i)}$. The areas are computed by taking the minimum or maximum value of the linear functions of the box coordinates.

5. Experiments

The SVOL dataset⁴ is curated upon the ImageNet-VID dataset [66] and three different sketch datasets with varying levels of abstraction: **Sketchy** [20] (least abstract), **TU-Berlin** [34], and **QuickDraw** [69] (most abstract) (see Fig. 2).

5.1. Implementation Details

We uniformly sampled 32 frames from a video ($L = 32$), scaled them to 224×224 dimensions, and used them as an input $\mathcal{V} \in \mathbb{R}^{32 \times 3 \times 224 \times 224}$ (3 for RGB channels). Likewise, a sketch is rescaled to 224×224 size, and used as an input $\mathcal{S} \in \mathbb{R}^{224 \times 224}$. The number of CMT layers is set to two (*i.e.*, $l = 2$), and we used 10 object tokens per frame ($M = 10$), a total of 320 object tokens ($N = 320$). We adopt ResNet-50 and ResNet-18 [29] pre-trained on ImageNet [66] as our video and sketch backbone, respectively.

Due to excessive number of video-sketch pairs, we used an iteration-based batch sampler and randomly sampled a subset from all possible pairings for training. SVANet is trained using AdamW optimizer [51] with an initial learning rate of 10^{-4} and weight decay of 10^{-4} for a batch size of 16. The overall loss weights $\lambda_{L1} : \lambda_{giou} : \lambda_{cls}$ were set to 5 : 1 : 2 throughout training. We set different learning schedules (number of iterations and learning rate decay steps) for each sketch dataset as below since their sizes are different.

Settings	Sketchy	TU-Berlin	QuickDraw
# pairs (train)	1,545,801	215,040	2,958,400
Iterations	50,000	20,000	100,000
LR decay step*	30,000	6,000	30,000

*LR is linearly decayed by a factor of 10 at every LR decay step.

5.2. Experimental Setups

Evaluation metrics. We adopt two evaluation metrics for SVOL: 1) \mathbf{R}_μ^k denotes the percentage of samples that have at least one correct result in top- k retrieved results, *i.e.*, Recall@ k , where the correct results indicate that IoU with ground truth is larger than the threshold μ . (we specifically use $k = 1, 5$ and $\mu = 0.5, 0.7$); 2) **mIoU** averages the IoU

⁴Details about datasets, curation process, and statistics are in Appendix A.

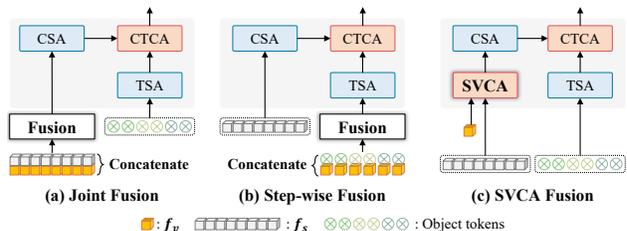


Figure 4. **Three instantiations of sketch-video fusion** contextualize video and sketch information in different ways: **(a) joint fusion:** copy f_s by the size of f_v , concatenate, and fuse them via MLP. **(b) step-wise fusion:** copy f_s by the number of object tokens, concatenate, and fuse them via MLP. The sketch and video representations are later fused through CTCA. **(c) SVCA fusion (ours):** fuse f_s and f_v with SVCA.

between predicted boxes and ground truth boxes over all testing samples to compare the overall performance.

Baselines. We set image-level sketch object localization approaches [65, 74] as the SVOL baselines. We find significant room for improvement as they were designed to be conditioned on a single frame rather than an entire video sequence. In addition, we present several instantiations of sketch-video fusion on SVANet, as shown in Fig. 4, and compare them with our final model.

5.3. Comparative Study

We benchmark the model performance on the SVOL task using three different sketch datasets. The results are shown in Table 1. Since image-level baselines (CMA, Sketch-DETR) make predictions at each frame, they neglect the global video context. In contrast, SVANet not only considers spatial context but also effectively models temporal information of video, thereby outperforming them by a significant margin in all metrics across three sketch datasets. Especially, SVANet improves mIoU by 29.4%, 17.7%, and 16.8% over Sketch-DETR on **Sketchy**, **TU-Berlin**, and **QuickDraw**. Also, our final SVANet yields the best results among the several model variants (joint, step-wise), implying the effectiveness of our attention-based fusion. Moreover, we use all three sketch datasets as a single set of query sketches (denoted as ALL in Table 1) to see how the model performs when the same category contains sketches of different styles. Overall, model performances are diminished as a result of a greater diversity of sketch samples. However, SVANet shows only 0.25%p mIoU drop compared to the results on **TU-Berlin**, which means that SVANet is quite robust to sketch style variations. Lastly, we compare two backbones for video encoding: 3D CNN (S3D [85]) vs. 2D CNN (ResNet50 [29]). We expected the more sophisticated 3D CNN to work better, but 2D CNN outperformed 3D CNN. This shows that CMT can sufficiently complement the temporal modeling capability of 3D CNN.

Method (backbone)	Sketchy					TU-Berlin					QuickDraw					ALL (S U T U Q)				
	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU
CMA [†] [74] (R50)	23.18	14.89	39.76	21.29	19.76	20.25	14.55	36.87	20.80	18.64	22.69	15.53	39.86	20.80	21.52	18.89	11.41	33.53	17.23	16.14
Sketch-DETR [†] [65] (R50)	28.78	18.56	46.65	26.50	26.09	30.75	18.97	47.76	27.54	26.24	31.10	19.47	49.39	31.05	28.59	28.23	16.74	44.21	25.54	24.30
SVANet/joint (R50)	33.86	22.56	52.84	30.57	31.46	31.14	19.48	50.17	28.21	29.38	33.95	20.12	54.77	34.11	31.89	29.53	16.50	47.90	27.18	28.59
SVANet/step-wise (R50)	33.31	22.81	53.00	31.07	30.29	30.23	18.19	50.66	27.20	30.41	32.05	21.17	56.34	35.39	31.98	29.64	17.61	48.21	26.29	27.99
SVANet (S3D [85])	32.89	21.11	48.08	27.07	30.83	29.43	18.25	45.72	24.34	28.00	30.86	19.24	46.80	25.48	29.37	27.88	17.26	43.33	22.86	27.87
SVANet (R50 [29])	35.60	23.19	54.06	32.95	33.76	32.10	19.60	51.61	30.94	30.89	34.47	22.30	58.13	37.88	33.40	31.80	18.52	51.44	29.90	30.64

Table 1. **Comparison of SVANet with baselines.** SVANet significantly outperforms baselines on three sketch datasets and on combined dataset (ALL), where we use only overlapping categories between three datasets. [†] indicates the re-implementation based on our settings.

Method	Sketchy→TU-Berlin					Sketchy→QuickDraw					Method	Seen →Unseen categories				
	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU		R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU
CMA [†] [74]	38.03	26.20	49.70	39.07	30.97	40.06	29.89	49.83	37.69	32.45	CMA [†] [74]	18.16	6.89	25.52	9.78	13.37
Sketch-DETR [†] [65]	43.49	32.25	51.71	44.39	36.76	46.02	37.87	59.50	45.34	40.21	Sketch-DETR [†] [65]	25.13	13.20	34.87	18.45	22.51
SVANet (Ours)	54.74	46.67	69.90	56.57	49.01	55.03	47.56	72.04	58.87	49.74	SVANet (Ours)	30.13	18.58	41.18	25.51	29.89

(a) The models are trained with Sketchy dataset and evaluated on QuickDraw or TU-Berlin dataset. To solely see the effect of sketch style differences, we use the same video samples in both training and evaluation, and overlapping categories between the two sketch datasets.

(b) The models are trained on 14 categories of the Sketchy dataset and evaluated on the remaining 5 categories: aircraft, bear, cat, cow, and dog.

Table 2. **Transfer evaluation on (a) unseen datasets and (b) unseen categories.**

5.4. Transfer Evaluation

The prediction space of our SVANet is not limited to a fixed set of categories. By design, it is possible to match even an unseen sketch to the most similar object by comparing feature-level similarity. For SVOL system to be more practical in real-world applications, they should be able to operate well even with sketches of various shapes and styles. In addition, there should be no constraint that operate only for limited categories, such as object detectors. To this end, we devise two transfer tasks to evaluate the generalization capability of the SVOL systems in two aspects: (i) dataset-level transfer and (ii) category-level transfer.

Formally, we define the transfer evaluation setup as follows. Let \mathcal{V} , $\{\mathcal{S}_A, \mathcal{S}_B\}$, and $\{\mathcal{C}_A, \mathcal{C}_B\}$ be a video dataset, sketch datasets, and sets of categories in which \mathcal{S}_A and \mathcal{S}_B overlap with \mathcal{V} , respectively. For dataset-level transfer task, we train the SVOL model on \mathcal{V} and \mathcal{S}_A , and evaluate on \mathcal{V} and \mathcal{S}_B , only for categories $\mathcal{C}_A \cap \mathcal{C}_B$. For category-level transfer task, we first split a sketch dataset \mathcal{S}_A into two subsets: \mathcal{S}_A^1 and \mathcal{S}_A^2 , where they are mutually exclusive w.r.t. categories, *i.e.*, $\mathcal{C}_A^1 \cap \mathcal{C}_A^2 = \emptyset$. Then, we train the SVOL model on \mathcal{V} and \mathcal{S}_A^1 , and evaluate on \mathcal{V} and \mathcal{S}_A^2 . We note that there can be more variations to evaluate the transferability of the SVOL system.

Transfer to unseen dataset. We study the transferability of the SVOL models across the sketch datasets with style differences (*e.g.*, line thickness, abstraction degree, *etc.*). The models are trained with Sketchy dataset and evaluated on QuickDraw or TU-Berlin dataset. To solely examine the transferability on unseen datasets, we use the same video samples in both training and evaluation, and overlapping categories between the two sketch datasets. The results are shown in Table 2a. The overall transfer performances across the datasets is much higher than the performance of mod-

els that are solely trained on dataset itself (Table 1), since transfer settings use the same video set as in training. We observe that SVANet significantly outperforms baselines in dataset-level transfer, indicating that it effectively learns class-discriminative features independent of sketch style differences. Meanwhile, we expected transfer to TU-Berlin to show better results than transfer to QuickDraw as TU-Berlin appears to be closer to Sketchy than QuickDraw in terms of visual similarity. Contrary to our expectation, transfer to QuickDraw shows better results than transfer to TU-Berlin. We understand this is because the system constructs categorical embedding space by matching the query sketch and video objects based on the key features (*e.g.*, cat’s whiskers, rabbit’s ears, *etc.*), rather than merely comparing their overall visual appearance. The results suggest that implicit similarity such as symbolic representation of sketches are more important for accurate object localization than explicit similarity such as line thickness or proportion.

Transfer to unseen categories. In Table 2b, we evaluate transferability of SVOL models at the category-level. We use 14 categories of the Sketchy dataset for training, and the remaining 5 categories (aircraft, bear, cat, cow, dog) for evaluation. Compared to the results in Table 1, we observe that SVANet degrades 3.87%p in mIoU since it has never learned which video objects to match the query sketch with. Despite this, SVANet outperforms the baselines when evaluated on unseen categories, implying that SVANet has learned more generalizable representations that can reason about the implicit similarities between sketches and video objects. This enables SVANet to closely embed sketches of the same category in the feature space.

def. CSA TSA	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	layers	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	pfsm	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU
✓	32.98	19.76	51.74	31.19	30.72	1	33.53	19.95	47.05	24.17	32.30	✗	18.36	6.78	33.54	13.64	22.34
✓ ✓	34.39	21.85	53.69	32.28	32.52	2	35.60	23.19	54.06	32.95	33.76	✓	35.60	23.19	54.06	32.95	33.76
✓ ✓	33.83	20.89	52.29	31.73	31.69	3	35.14	20.18	56.29	32.99	32.77	△	+17.24	+16.41	+20.52	+19.31	+11.42
✓ ✓ ✓	35.60	23.19	54.06	32.95	33.76	4	35.20	22.77	58.34	37.39	33.28		△: performance gain.				

(a) CMT attention operations. default: SVCA + CTCA.

(b) CMT depth.

(c) Per-frame set matching (pfsm).

tokens	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	obj	ℓ ₁	gIoU	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU	frames	R _{0.5} ¹	R _{0.7} ¹	R _{0.5} ⁵	R _{0.7} ⁵	mIoU
5	34.53	20.50	47.61	25.21	32.34	✓	✓		33.64	20.75	46.40	24.37	32.44	16	32.66	18.79	54.24	31.85	31.25
10	35.60	23.19	54.06	32.95	33.76	✓		✓	32.13	18.16	42.32	23.21	31.72	32	35.60	23.19	54.06	32.95	33.76
15	33.89	22.98	54.98	32.64	32.69	✓	✓	✓	35.60	23.19	54.06	32.95	33.76	64	34.91	22.83	53.72	31.23	32.48

(d) Object tokens per frame.

(e) Loss. obj loss is set as default as it is essential.

(f) Input video density.

Table 3. Ablative experiments. Our settings are marked in gray. All experiments are conducted on the Sketchy dataset.

5.5. Ablative Study

CMT attention operations. We study the effect of four attention operations of CMT in Table 3a. Here, SVCA and CTCA are set as default since they are indispensable for making predictions in our design. Each is responsible for modeling interaction between sketch and video, and transforming object tokens into predictions conditioned on the sketch-video joint representations. CSA models the global context of the input sequence and TSA models relationships between object tokens. The default setting work fairly well (mIoU=30.72%), yet SVANet shows better performance with the addition of CSA (+2.33%p) or TSA (+0.97%p). In particular, CSA plays a crucial role in object localization in video since it is in charge of temporal modeling, thus leading to a substantial performance increase. We confirm that all CMT components operate collaboratively on the SVOL task, as they achieve the best performance when used together.

CMT depth. We examine the effect of varying the CMT depth (*i.e.*, number of layers) in Table 3b. A single layer of CMT does not provide sufficient contextualization, resulting in poor R⁵ performance. The overall performance seems balanced between two to four layers. For R⁵ metric, the deeper the layer, the better the performance, and the best performance is achieved with four CMT layers. However, for more strict R¹ and mIoU metrics, two layers perform the best. Therefore, we made two layers as our default setting.

Per-frame set matching. A straightforward way for training SVANet is to match all predictions with all ground truths as a whole. Although simple, it requires learning all N object tokens simultaneously, regardless of frame order. On the contrary, our per-frame set matching strategy divides N object tokens into L subsets of M object tokens, then matches only a subset to ground truths of its corresponding frame. Although set matching is performed frame-by-frame, SVANet can still make predictions in parallel. We compare our strategy to the straightforward approach in Table 3c. Overall, using per-frame set matching resulted in a significant performance improvement. We see this is because our strategy not only eases optimization by reducing the set matching

complexity, but also brings a strong positional inductive bias for object tokens (see empirical evidence in Fig. 6c).

Number of object tokens. In order to see the effect of the number of object tokens used in the CMT layers, we varied their number in Table 3d. Too few tokens (=5) limit sufficient interactions between foregrounds and backgrounds (*i.e.*, No Object), resulting in poor performance, especially for R⁵ metric. On the other hand, too many tokens (=15) diminish performance by producing unnecessary backgrounds. Having 10 object tokens per frame provides a good balance between foreground and background, resulting in a good performance. As we utilize a per-frame set matching strategy, we set the number of object tokens per frame to 10 for the entire 32 frames, thus using a total of 320 object tokens.

Loss components. In Table 3e, we toggle the loss components on and off to understand their impact on training. The objectness loss is used in all cases since it is essential to determine whether a prediction contains the target object. When either ℓ₁ or gIoU [64] loss is disabled, performances drop drastically, especially in R⁵ metric. This implies that both losses are not only important for accurate box localization, but also for performing overall predictions well. As we obtain the best results when using all three losses, we confirm that scale-sensitive ℓ₁ and scale-invariant gIoU losses operate complementarily with each other.

Sampling density of video frames. In Table 3f, we study the effect of frame sampling density on input video. We uniformly sample a fixed number of frames across the video and use them as an input to SVANet. By default, we use 32 frames. Compared to the baseline, 16 frames show a particularly sharp performance drop on the strictest metric R_{0.7}¹, and 64 frames show overall sub-optimal performance. This is because sparse sampling enables faster processing with less memory, but can easily miss important details since it provides less information. In contrast, dense sampling provides more information, but if the motion of objects is not large, it can be redundant and rather hinder optimization. Here, we study only simple uniform sampling, but different means of sampling may achieve different results.

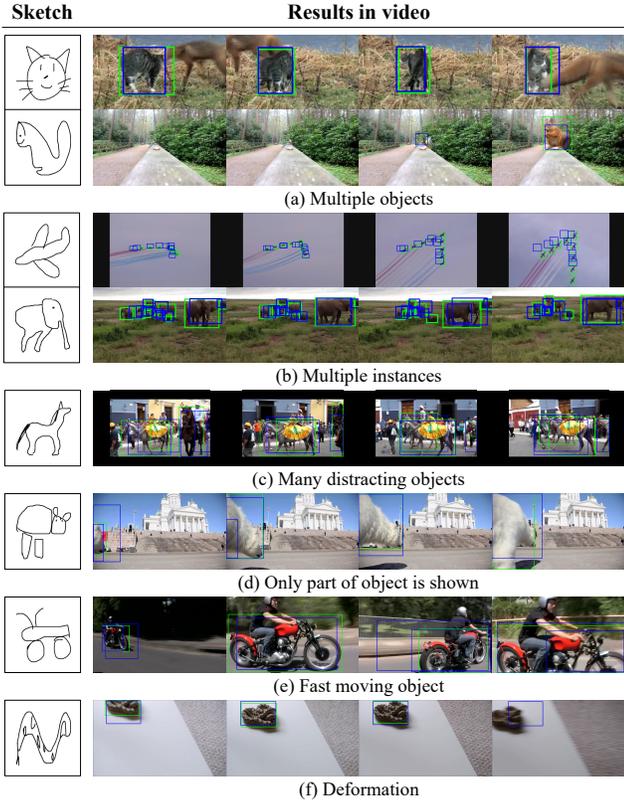


Figure 5. **Qualitative results** of SVANet on QuickDraw dataset. Green and blue boxes represent ground truths and predictions, respectively. SVANet performs well in various challenging scenarios, including: (a) when there are confusable objects; (b) multiple object instances appear in a video; (c) there are many distracting objects; (d) only part of the object is visible; (e) the target object moves quickly; (f) the appearance of the target object is not similar to query sketch.

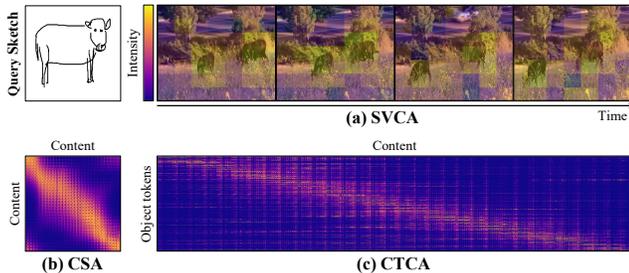


Figure 6. **Attention visualization.** The brighter (yellowish) the color, the higher the attention intensity. Given a query sketch, (a) SVCA mainly attends to regions where the target objects are located; (b) CSA gives a higher attention weight to temporally neighboring contents; (c) CTCA learns which content object tokens should mainly focus on temporally.

5.6. Qualitative Analysis

SVOL results. We present qualitative results of SVANet in Fig. 5 to illustrate how it works in practice.⁵ Our system successfully recognizes the objects that correspond to the

⁵We present additional qualitative results in Appendix B.

query sketch and accurately localizes their bounding boxes in a variety of challenging conditions. SVANet works well even when: (a) there are two confusable objects; (b) multiple object instances appear in a video; (c) there are lots of distracting objects; (d) only part of the object is appearing; (e) the target object moves quickly; (f) the appearance of the query sketch is not similar to that of the target object.

CMT attention visualization. In order to understand the behavior of CMT, we visualize its attention maps in Fig. 6. Our observations are as follows: (a) SVCA learns *where to look*, as such, the highlighted area on the attention map aligns well with the actual locations of the sketch object. (b) CSA learns deeper correlation between temporally adjacent sequences when modeling temporal context. (c) CTCA learns *when to look*, thereby giving temporal inductive bias for object tokens in conjunction with per-frame set matching.

6. Conclusion

We have introduced a new challenging task termed Sketch-b Video Object Localization (SVOL), which aims to localize objects in video that match a given query sketch. To tackle this task, we propose a strong baseline model called SVANet, which addresses the temporal context of video and the domain gap between sketches and videos. Our model, SVANet, utilizes two key designs to solve the SVOL task as a set prediction problem: a Cross-modal Transformer and per-frame set matching. To verify our approach, we conduct extensive experiments on a newly curated SVOL dataset and have found that SVANet significantly outperforms image-level approaches by 29.4%, 17.7%, and 16.8%. We further analyze the behavior of SVANet through comprehensive ablations and visualizations. Last but not least, we found that SVANet generalizes well to unseen datasets and novel categories, implying its scalability in real-world scenarios. We hope our work will inspire further research in the field of sketch-based object localization.⁶

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing Moments in Video With Natural Language. In *ICCV*, pages 5803–5812, 2017. 1, 4
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded Up Robust Features. In *ECCV*, pages 404–417, 2006. 4
- [3] Ayan Kumar Bhunia, Viswanatha Reddy Gajjala, Subhadeep Koley, Rohit Kundu, Aneeshan Sain, Tao Xiang, and Yi-Zhe Song. Doodle It Yourself: Class Incremental Learning by Drawing a Few Sketches. In *CVPR*, pages 2293–2302, 2022. 1, 3

⁶Limitations, future work, and broader impacts are in Appendix C.

- [4] Ayan Kumar Bhunia, Yongxin Yang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Sketch Less for More: On-The-Fly Fine-Grained Sketch-Based Image Retrieval. In *CVPR*, pages 9779–9788, 2020. [20](#)
- [5] Federico Boniardi, Abhinav Valada, Wolfram Burgard, and Gian Diego Tipaldi. Autonomous Indoor Robot Navigation Using a Sketch Interface for Drawing Maps and Routes. In *ICRA*, pages 2896–2901. IEEE, 2016. [3](#)
- [6] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multi-modal transformers. In *CVPR*, pages 4985–4995, 2022. [4](#)
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language Models Are Few-Shot Learners. In *NeurIPS*, pages 1877–1901, 2020. [4](#)
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-End Object Detection With Transformers. In *ECCV*, pages 213–229, 2020. [2, 4](#)
- [9] Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. WebQA: Multihop and Multimodal QA. In *CVPR*, pages 16495–16504, 2022. [4](#)
- [10] Wengling Chen and James Hays. Sketchygan: Towards Diverse and Realistic Sketch to Image Synthesis. In *CVPR*, pages 9416–9425, 2018. [1, 3](#)
- [11] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer Tracking. In *CVPR*, pages 8126–8135, 2021. [1, 4](#)
- [12] Yihong Chen, Yue Cao, Han Hu, and Liwei Wang. End-to-End Video Object Detection With Spatial-Temporal Transformers. In *CVPR*, pages 10337–10346, 2020. [3](#)
- [13] Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, and Rongrong Ji. Siamese Box Adaptive Network for Visual Tracking. In *CVPR*, pages 6668–6677, 2020. [4](#)
- [14] Navneet Dalal and Bill Triggs. Histograms of Oriented Gradients for Human Detection. In *CVPR*, pages 886–893, 2005. [4](#)
- [15] Chaorui Deng, Qi Wu, Qingyao Wu, Fuyuan Hu, Fan Lyu, and Mingkui Tan. Visual Grounding via Accumulated Attention. In *CVPR*, pages 7746–7755, 2018. [4](#)
- [16] Jiajun Deng, Zhengyuan Yang, Tianlang Chen, Wengang Zhou, and Houqiang Li. TransVG: End-to-End Visual Grounding With Transformers. In *ICCV*, pages 1769–1779, 2021. [4](#)
- [17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*, 2018. [4](#)
- [18] Sounak Dey, Pau Riba, Anjan Dutta, Josep Lladós, and Yi-Zhe Song. Doodle To Search: Practical Zero-Shot Sketch-Based Image Retrieval. In *CVPR*, pages 2179–2188, 2019. [3, 15](#)
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv preprint arXiv:2010.11929*, 2020. [4](#)
- [20] Mathias Eitz, James Hays, and Marc Alexa. How Do Humans Sketch Objects? *ACM TOG*, pages 1–10, 2012. [2, 8, 15, 16](#)
- [21] Hehe Fan and Yi Yang. Person Tube Retrieval via Language Description. In *AAAI*, pages 10754–10761, 2020. [4](#)
- [22] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-Shot Object Detection With Attention-RPN and Multi-Relation Detector. In *CVPR*, pages 4013–4022, 2020. [3](#)
- [23] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect To Track and Track To Detect. In *ICCV*, pages 3038–3046, 2017. [3](#)
- [24] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding. *arXiv preprint arXiv:1606.01847*, 2016. [4](#)
- [25] Yaroslav Ganin and Victor Lempitsky. Unsupervised Domain Adaptation by Backpropagation. In *ICML*, pages 1180–1189, 2015. [20](#)
- [26] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-Adversarial Training of Neural Networks. *JMLR*, pages 2096–2030, 2016. [20](#)
- [27] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. TALL: Temporal Activity Localization via Language Query. In *ICCV*, pages 5267–5275, 2017. [4](#)
- [28] Wei Han, Pooya Khorrami, Tom Le Paine, Prajit Ramachandran, Mohammad Babaeizadeh, Honghui Shi, Jianan Li, Shuicheng Yan, and Thomas S Huang. Seq-NMS for Video Object Detection. *arXiv preprint arXiv:1602.08465*, 2016. [3](#)
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, pages 770–778, 2016. [6, 8, 9](#)
- [30] Ting-I Hsieh, Yi-Chen Lo, Hwann-Tzong Chen, and Tyng-Luh Liu. One-Shot Object Detection With Co-Attention and Co-Excitation. In *NeurIPS*, 2019. [1, 3](#)
- [31] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural Language Object Retrieval. In *CVPR*, pages 4555–4564, 2016. [4](#)
- [32] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-Image Translation With Conditional Adversarial Networks. In *CVPR*, pages 1125–1134, 2017. [1, 3](#)
- [33] Ying Jiang, Congyi Zhang, Hongbo Fu, Alberto Cannavò, Fabrizio Lamberti, Henry YK Lau, and Wenping Wang. HandPainter-3D Sketching in VR with Hand-Based Physical Proxy. In *CHI*, pages 1–13, 2021. [3](#)
- [34] Jonas Jongejans, Henry Rowley, Takashi Kawashima, Jongmin Kim, and Nick Fox-Gieg. The Quick, Draw!-AI Experiment. *Mount View, CA, accessed Feb*, page 4, 2016. [2, 8, 15, 16](#)
- [35] Bumsoo Kim, Junhyun Lee, Jaewoo Kang, Eun-Sol Kim, and Hyunwoo J Kim. Hotr: End-to-End Human-Object Interaction Detection With Transformers. In *CVPR*, pages 74–83, 2021. [4](#)
- [36] Chen Kong, Dahua Lin, Mohit Bansal, Raquel Urtasun, and Sanja Fidler. What Are You Talking About? Text-to-Image Coreference. In *CVPR*, pages 3558–3565, 2014. [1, 4](#)
- [37] Harold W Kuhn. The Hungarian Method for the Assignment Problem. *Naval research logistics quarterly*, pages 83–97, 1955. [7](#)
- [38] Kin Chung Kwan and Hongbo Fu. Mobi3dsketch: 3D Sketching in Mobile AR. In *CHI*, pages 1–11, 2019. [1, 3](#)
- [39] Sumin Lee, Sangmin Woo, Yeonju Park, Muhammad Adi Nugroho, and Changick Kim. Modality Mixer for Multi-

- modal Action Recognition. In *WACV*, pages 3298–3307, 2023. 4
- [40] Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. TVQA+: Spatio-Temporal Grounding for Video Question Answering. *arXiv preprint arXiv:1904.11574*, 2019. 4
- [41] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of Siamese Visual Tracking With Very Deep Networks. In *CVPR*, pages 4282–4291, 2019. 4
- [42] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High Performance Visual Tracking With Siamese Region Proposal Network. In *CVPR*, pages 8971–8980, 2018. 1, 4
- [43] Naihuan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. Neural Speech Synthesis With Transformer Network. In *AAAI*, volume 33, pages 6706–6713, 2019. 4
- [44] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. OSCAR: Object-Semantics Aligned Pre-training for Vision-Language Tasks. In *ECCV*, pages 121–137. Springer, 2020. 4
- [45] Pengpeng Liang, Yifan Wu, Hu Lu, Liming Wang, Chunyuan Liao, and Haibin Ling. Planar object tracking in the wild: A benchmark. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 651–658. IEEE, 2018. 1
- [46] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 3
- [47] Ioana Literat. “a pencil for your thoughts”: Participatory drawing as a visual research method with children and youth. *International Journal of Qualitative Methods*, 12(1):84–98, 2013. 1
- [48] Fang Liu, Xiaoming Deng, Changqing Zou, Yu-Kun Lai, Keqi Chen, Ran Zuo, Cuixia Ma, Yong-Jin Liu, and Hongan Wang. Scenesketcher-V2: Fine-Grained Scene-Level Sketch-Based Image Retrieval Using Adaptive GCNs. *TIP*, 31:3737–3751, 2022. 3
- [49] Jialu Liu. Image Retrieval Based on Bag-of-Words Model. *arXiv preprint arXiv:1304.5168*, 2013. 1, 4
- [50] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single Shot Multibox Detector. In *ECCV*, pages 21–37, 2016. 3
- [51] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. *arXiv preprint arXiv:1711.05101*, 2017. 8
- [52] David G Lowe. Distinctive Image Features From Scale-Invariant Keypoints. *IJCV*, pages 91–110, 2004. 4
- [53] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViL-BERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. *NeurIPS*, 32, 2019. 4
- [54] Zhaoliang Lun, Matheus Gadelha, Evangelos Kalogerakis, Subhransu Maji, and Rui Wang. 3D Shape Reconstruction From Sketches via Multi-View Convolutional Networks. In *3DV*, pages 67–77. IEEE, 2017. 3
- [55] Anton Osokin, Denis Sumin, and Vasily Lomakin. OS2D: One-Stage One-Shot Object Detection by Matching Anchor Features. In *ECCV*, pages 635–652, 2020. 3
- [56] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment Matching for Multi-Source Domain Adaptation. In *ICCV*, pages 1406–1415, 2019. 20
- [57] Tiziano Portenier, Qiyang Hu, Attila Szabo, Siavash Arjomand Bigdeli, Paolo Favaro, and Matthias Zwicker. Faceshop: Deep Sketch-Based Face Image Editing. *arXiv preprint arXiv:1804.08972*, 2018. 1, 3
- [58] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, pages 8748–8763, 2021. 4
- [59] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving Language Understanding by Generative Pre-Training. 2018. 4
- [60] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language Models Are Unsupervised Multitask Learners. *OpenAI blog*, page 9, 2019. 4
- [61] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-Shot Text-To-Image Generation. In *ICML*, pages 8821–8831. PMLR, 2021. 4
- [62] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. In *CVPR*, pages 779–788, 2016. 3
- [63] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection With Region Proposal Networks. In *NeurIPS*, 2015. 3, 4
- [64] Hamid Rezaatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *CVPR*, pages 658–666, 2019. 7, 10
- [65] Pau Riba, Sounak Dey, Ali Furkan Biten, and Josep Lladós. Localizing Infinity-Shaped Fishes: Sketch-Guided Object Localization in the Wild. *arXiv preprint arXiv:2109.11874*, 2021. 1, 2, 4, 8, 9, 17
- [66] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet Large Scale Visual Recognition Challenge. *IJCV*, pages 211–252, 2015. 2, 8, 15, 16
- [67] Aneeshan Sain, Ayan Kumar Bhunia, Yongxin Yang, Tao Xiang, and Yi-Zhe Song. Stylemeup: Towards Style-Agnostic Sketch-Based Image Retrieval. In *CVPR*, pages 8504–8513, 2021. 2, 3, 15
- [68] Daisuke Sakamoto, Koichiro Honda, Masahiko Inami, and Takeo Igarashi. Sketch and Run: A Stroke-Based Interface for Home Robots. In *CHI*, pages 197–200, 2009. 3
- [69] Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. The Sketchy Database: Learning To Retrieve Badly Drawn Bunnies. *ACM TOG*, pages 1–12, 2016. 2, 8, 15, 16
- [70] Jifei Song, Qian Yu, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. Deep Spatial-Semantic Attention for Fine-Grained Sketch-Based Image Retrieval. In *ICCV*, pages 5551–5560, 2017. 3
- [71] Rui Su, Qian Yu, and Dong Xu. Styvgbert: A Visual-Linguistic Transformer Based Framework for Spatio-Temporal Video Grounding. In *ICCV*, pages 1533–1542, 2021. 1, 4

- [72] Haifeng Sun, Jiaqing Xu, Jingyu Wang, Qi Qi, Ce Ge, and Jianxin Liao. DLI-Net: Dual Local Interaction Network for Fine-Grained Sketch-Based Image Retrieval. *TCSVT*, 32(10):7177–7189, 2022. [3](#)
- [73] Paul Taelle, Rachel Blagojevic, Tracy Hammond, Samantha Ray, Josh Cherian, and Jung In Koh. Sketchrec 2023: 3rd workshop on sketch recognition. In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces*, pages 1–1, 2023. [1](#)
- [74] Aditay Tripathi, Rajath R Dani, Anand Mishra, and Anirban Chakraborty. Sketch-Guided Object Localization in Natural Images. In *ECCV*, pages 532–547, 2020. [1](#), [4](#), [8](#), [9](#)
- [75] Laurens Van der Maaten and Geoffrey Hinton. Visualizing Data Using T-SNE. *JMLR*, 2008. [18](#)
- [76] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All You Need. In *NeurIPS*, pages 5998–6008, 2017. [2](#), [4](#), [5](#)
- [77] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching Networks for One Shot Learning. In *NeurIPS*, 2016. [3](#)
- [78] Fang Wang, Le Kang, and Yi Li. Sketch-Based 3D Shape Retrieval Using Convolutional Neural Networks. In *CVPR*, pages 1875–1883, 2015. [3](#)
- [79] Han Wang, Jun Tang, Xiaodong Liu, Shanyan Guan, Rong Xie, and Li Song. Ptseformer: Progressive temporal-spatial enhanced transformer towards video object detection. In *ECCV*, pages 732–747. Springer, 2022. [3](#)
- [80] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-DeepLab: End-to-End Panoptic Segmentation With Mask Transformers. In *CVPR*, pages 5463–5474, 2021. [4](#)
- [81] Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. Sketch Your Own Gan. In *ICCV*, pages 14050–14060, 2021. [1](#), [3](#)
- [82] Sangmin Woo, Sumin Lee, Yeonju Park, Muhammad Adi Nugroho, and Changick Kim. Towards Good Practices for Missing Modality Robust Action Recognition. *arXiv preprint arXiv:2211.13916*, 2022. [4](#)
- [83] Sangmin Woo, Jinyoung Park, Inyong Koo, Sumin Lee, Minki Jeong, and Changick Kim. Explore and Match: End-to-End Video Grounding With Transformer. *arXiv preprint arXiv:2201.10168*, 2022. [1](#), [4](#)
- [84] Haiping Wu, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Sequence Level Semantics Aggregation for Video Object Detection. In *ICCV*, pages 9217–9225, 2019. [3](#)
- [85] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking Spatiotemporal Feature Learning: Speed-Accuracy Trade-Offs in Video Classification. In *ECCV*, pages 305–321, 2018. [8](#), [9](#)
- [86] Peng Xu, Timothy M Hospedales, Qiyue Yin, Yi-Zhe Song, Tao Xiang, and Liang Wang. Deep Learning for Free-Hand Sketch: A Survey. *TPAMI*, 2022. [1](#), [3](#)
- [87] Peng Xu, Kun Liu, Tao Xiang, Timothy M Hospedales, Zhanyu Ma, Jun Guo, and Yi-Zhe Song. Fine-Grained Instance-Level Sketch-Based Video Retrieval. *TCSVT*, pages 1995–2007, 2020. [3](#)
- [88] Shuai Yang, Zhangyang Wang, Jiaying Liu, and Zongming Guo. Deep Plastic Surgery: Robust and Controllable Image Editing With Human-Drawn Sketches. In *ECCV*, pages 601–617. Springer, 2020. [3](#)
- [89] Zhao Yang, Jiaqi Wang, Yansong Tang, Kai Chen, Hengshuang Zhao, and Philip HS Torr. Lavt: Language-aware vision transformer for referring image segmentation. In *CVPR*, pages 18155–18165, 2022. [1](#), [4](#)
- [90] Qian Yu, Feng Liu, Yi-Zhe Song, Tao Xiang, Timothy M Hospedales, and Chen-Change Loy. Sketch Me That Shoe. In *CVPR*, pages 799–807, 2016. [1](#), [3](#), [15](#)
- [91] Qian Yu, Jifei Song, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. Fine-Grained Instance-Level Sketch-Based Image Retrieval. *IJCV*, 129(2):484–500, 2021. [3](#), [20](#)
- [92] Hao Zhang, Aixin Sun, Wei Jing, Guoshun Nan, Liangli Zhen, Joey Tianyi Zhou, and Rick Siow Mong Goh. Video Corpus Moment Retrieval With Contrastive Learning. In *SIGIR*, pages 685–695, 2021. [20](#)
- [93] Zhu Zhang, Zhou Zhao, Yang Zhao, Qi Wang, Huasheng Liu, and Lianli Gao. Where Does It Exist: Spatio-Temporal Video Grounding for Multi-Form Sentences. In *CVPR*, pages 10668–10677, 2020. [4](#)
- [94] Qianyu Zhou, Xiangtai Li, Lu He, Yibo Yang, Guangliang Cheng, Yunhai Tong, Lizhuang Ma, and Dacheng Tao. Transvod: end-to-end video object detection with spatial-temporal transformers. *TPAMI*, 2022. [3](#)

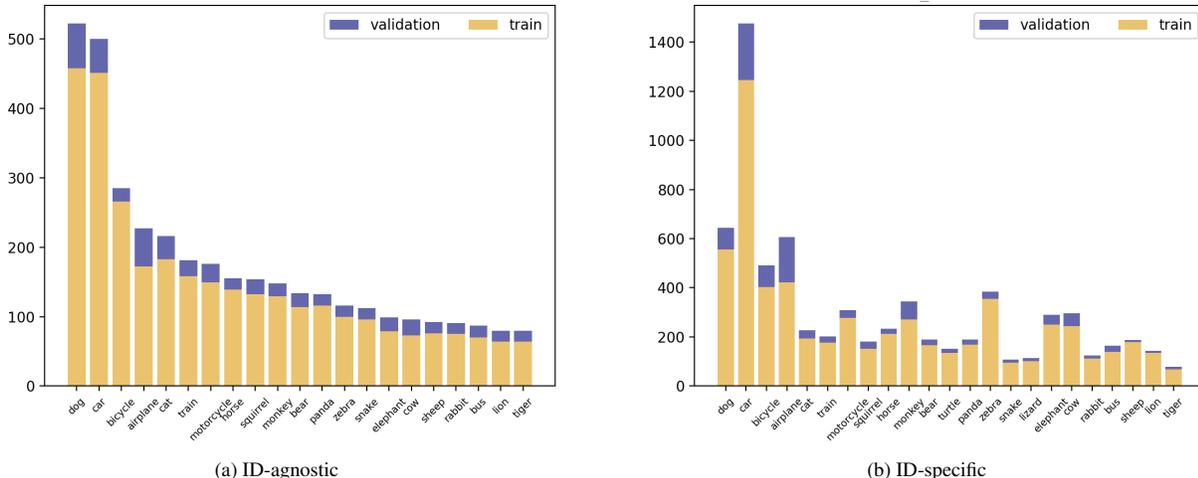


Figure 7. **Video dataset [66] class distribution** when objects of the same category are (a) counted as a whole (*i.e.*, *agnostic* to instance ID) or (b) counted individually (*i.e.*, *sensitive* to instance ID). The x-axis denotes the category and y-axis denotes the frequency.

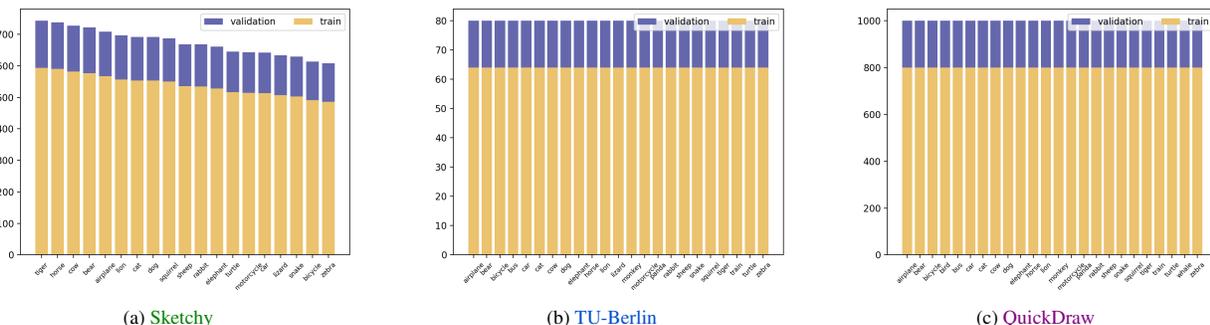


Figure 8. **Sketch dataset class distribution.** (a) *Sketchy* [69], (b) *TU-Berlin* [20], and (c) *QuickDraw* [34]. The x-axis denotes the class and y-axis denotes the frequency.

A. SVOL Dataset & Analysis

The SVOL dataset is built on multiple datasets [20, 34, 66, 69]. We consider only the categories that intersects between the video dataset [66] and the sketch datasets [20, 34, 69].

A.1. SVOL Dataset

ImageNet-VID [66] is built for video object detection task. It contains 5,354 snippets (train/val/test split is 3,852/555/937) that are annotated with 30 object categories, including vehicles (*e.g.*, airplane, bus, *etc.*) and animals (*e.g.*, bird, dog, *etc.*). Each object instance is annotated in the form of {video name, frame number, class label, instance id, bounding box}. We use a validation set for evaluation since test annotations are not publicly available.

Sketchy [69] is a large-scale collection of sketch-photo pairs that has 75,471 sketches belonging to 125 categories. Drawers are not allowed to directly trace objects; rather, sketches are drawn after seeing specific photographic objects. This forces the drawers to sketch from memory in the same way that a user of sketch-based image retrieval sys-

tems [18, 67, 90] would draw from a mental image of the desired object. 23 categories overlap with ImageNet-VID: airplane, bear, bicycle, car, cat, cow, dog, elephant, horse, lion, lizard, motorcycle, rabbit, sheep, snake, squirrel, tiger, turtle, zebra.

TU-Berlin [20] is a crowd-sourced sketch dataset composed of 20,000 unique sketches with 250 categories. The sketches are uniformly distributed over 250 object categories, which exhaustively cover the vast majority of objects seen in daily life. The median drawing time for each sketch is 86 seconds. Due to the low quality of some sketches, humans correctly identify just 73% of these hand-drawings. 21 categories overlap with ImageNet-VID: airplane, bear, bicycle, bus, car, cat, cow, dog, elephant, horse, lion, monkey, motorcycle, panda, rabbit, sheep, snake, squirrel, tiger, train, zebra.

QuickDraw [34] is a huge collection of 50 million sketches organized into 345 categories. Over 15 million players have contributed millions of sketches playing “Quick, Draw!” game⁷, where a neural network tries to guess the sketches.

⁷<https://quickdraw.withgoogle.com/>

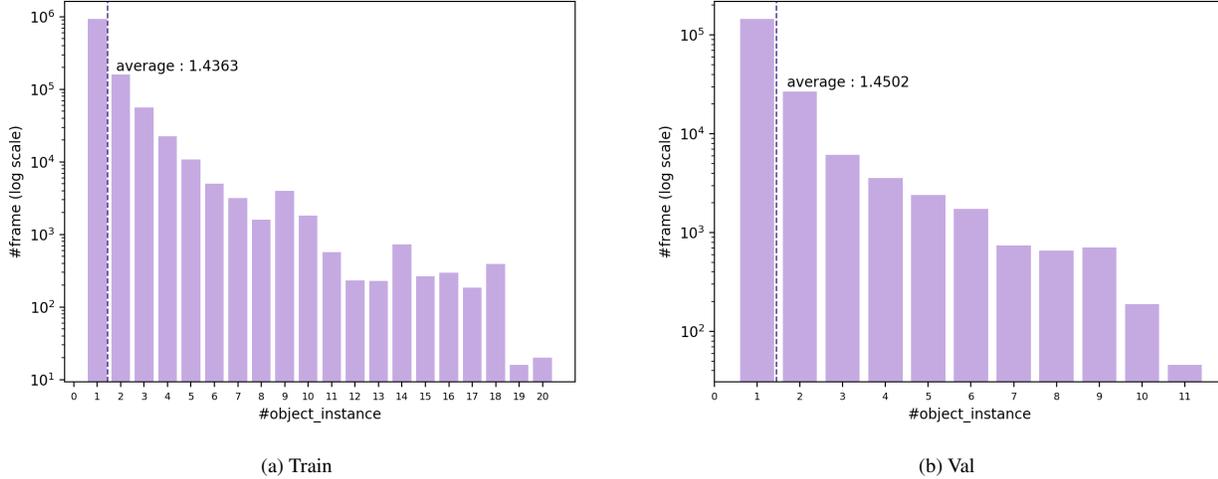


Figure 9. **Number of object instances per frame** in ImageNet-VID [66] (a) train and (b) validation split. The x-axis denotes the number of object instances and y-axis denotes the number of frames in a log scale.

The players are asked to draw a sketch of a given category in 20 seconds while the computer attempts to classify them. The way the sketches are collected results in a high degree of variety in the dataset, although most sketches are of low quality due to time limit. 24 categories overlap with ImageNet-VID.

Video \cap Sketch	Categories
ImageNet-VID \cap Sketchy (19 classes)	airplane, bear, bicycle, car, cat, cow, dog, elephant, horse, lion, lizard, motorcycle, rabbit, sheep, snake, squirrel, tiger, turtle, zebra
ImageNet-VID \cap TU-Berlin (21 classes)	airplane, bear, bicycle, bus, car, cat, cow, dog, elephant, horse, lion, monkey, motorcycle, panda, rabbit, sheep, snake, squirrel, tiger, train, zebra
ImageNet-VID \cap QuickDraw (24 classes)	airplane, bear, bicycle, bird, bus, car, cat, cow, dog, elephant, horse, lion, monkey, motorcycle, panda, rabbit, sheep, snake, squirrel, tiger, train, turtle, whale, zebra

A.2. SVOL Data Analysis

A.2.1 Frame Length Distribution

We use the video dataset from the ImageNet-VID dataset [66]. The train split has 3,862 videos that are fully annotated with the 30 object categories, yielding 866,870 bounding boxes for 1,122,397 frames. In validation split, 555 videos are fully annotated with the 30 object categories, resulting in 135,949 bounding box annotations for 176,126 frames. We summarize the statistics for the frame length distribution of ImageNet-VID dataset below.

Dataset	Split	min	max	mean	median
ImageNet-VID	Train	6	5492	290.6	180
	Val	11	2898	317.3	232

A.2.2 Class Histogram

Video dataset. We show the class histogram of ImageNet-VID [66] dataset in Fig. 7. Here, *ID-specific* refers to taking into account the identity (ID) of an object instance when counting the number, whereas *ID-agnostic* refers to not taking it into account. In *ID-agnostic*, the number is counted only once even if multiple object instances belonging to the same category appear in a video. For example, in the case of “car”, the number is around 500 without considering the track-id, but exceeds 1,400 with considering the track-id. This indicates that there are many scenes in the video in which multiple “car” object instances appear concurrently. We count only the object categories that are common in both the video and sketch datasets. The statistics for the class distribution of the SVOL video dataset are summarized below.

Dataset	Track-ID	Split	min	max	mean	median
ImageNet-VID	id-specific	Train	67	1246	276.8	194
		Val	9	229	47.7	29
	id-agnostic	Train	56	458	151.8	118
		Val	4	64	21.8	19

Sketch dataset. We show the class histograms of **Sketchy** [69] and **TU-Berlin** [20] datasets in Fig. 8. The **QuickDraw** [34] dataset has 1,000 sketch images per class are uniformly distributed and all train/val splits are 800/200. The number of sketches per class is relatively evenly distributed in the **Sketchy** dataset, however there is an imbalance between classes in the **TU-Berlin** dataset. The following table summarizes the statistics for the class distribution of SVOL sketch datasets.

Dataset	Split	min	max	mean	median
Sketchy	Train	486	594	539.5	535
	Val	122	149	135.4	134
TU-Berlin	Train	64	458	150.7	116
	Val	16	64	24.7	19
QuickDraw	Train	800	800	800	800
	Val	200	200	200	200

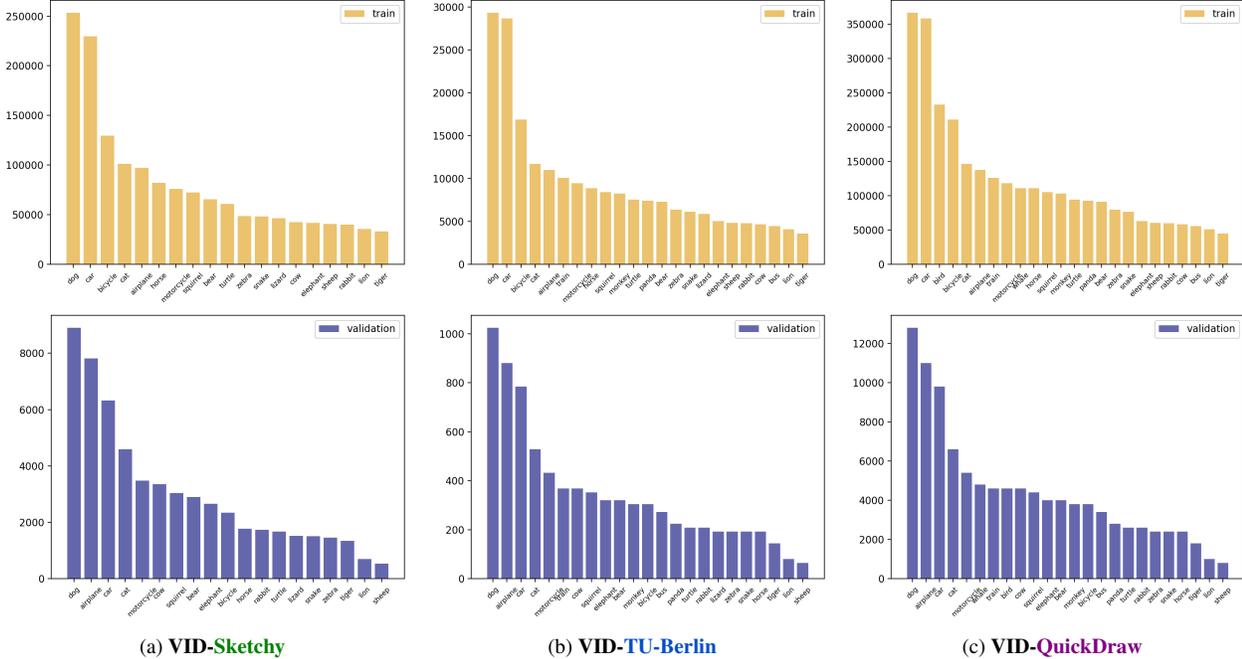


Figure 10. **Class distribution of video-sketch pair.** (a) VID-Sketchy, (b) VID-TU-Berlin, and (c) VID-QuickDraw. The x-axis denotes the class and y-axis denotes the frequency.

Number of object instances. Fig. 9 shows the distribution of the number of object instances per frame in the SVOL video dataset. The average object instances per video is 1.4363 in the train split and 1.4502 in the validation split.

A.3. Data Curation

The SVOL dataset is made up of a combination of video dataset and sketch datasets. To ensure that the SVOL evaluation set remain unseen in the training phase, we split videos and sketches into training and evaluation sets, respectively, and then construct a training set of SVOL with a combination of both training sets, and an evaluation set of SVOL with a combination of both evaluation sets. This guarantees the models to be evaluated on video-sketch pairs that are totally unseen throughout the training phase. While this setting is most closest to the actual environment in which the SVOL system operates, it requires the model to be generalized to both videos and sketches.

Formally, let $\{\mathcal{V}_{tr}, \mathcal{V}_{ev}\}$ train/eval video datasets, $\{\mathcal{S}_{tr}, \mathcal{S}_{ev}\}$ train/eval sketch datasets, and $\{\mathcal{C}_V, \mathcal{C}_S\}$ video/sketch category sets. For all categories that are common for video and sketch datasets, *i.e.*, $\forall c \in \mathcal{C}_V \cap \mathcal{C}_S$, we construct SVOL train set by pairing \mathcal{V}_{tr} and \mathcal{S}_{tr} , and SVOL eval set with \mathcal{V}_{ev} and \mathcal{S}_{ev} . Only video and sketch are paired when they have the same class label. The number of pairs generated for each video-sketch datasets are summarized in the table below (Fig. 10 shows the per-category distribution).

Split	VID-Sketchy	VID-TU-Berlin	VID-QuickDraw
Train	1,545,801	215,040	2,958,400
Eval	57,660	7,952	10,6400

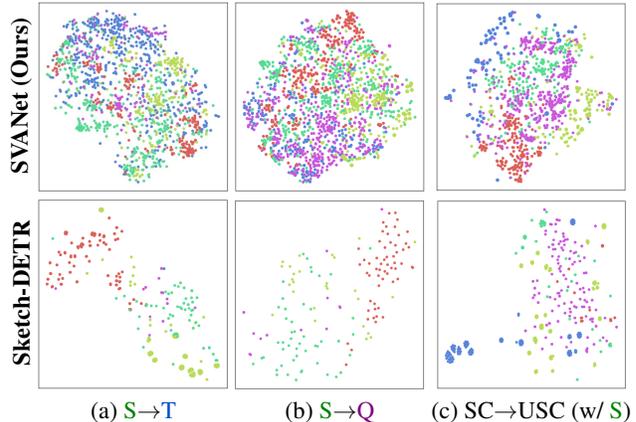


Figure 11. **Feature distribution of SVANet vs. Sketch-DETR [65]** when transfer is performed for the cases of (a) Sketchy \rightarrow TU-Berlin, (b) Sketchy \rightarrow QuickDraw, and (c) Seen \rightarrow Unseen Categories with the Sketchy dataset. Each data point represents the last hidden state of the CMT, and the color indicates the category it belongs to. We plot samples of 5 random categories with a confidence score higher than 0.9.

In practice, we only use videos that contain at least one query sketch object within 32 frames uniformly sampled from the video, and bounding box annotations that correspond to the sketch object are regarded as the ground truths for that pairing. A video can be paired more than once since it can contain multiple objects. We note that the class label is only a means for pairing and is not considered in actual training.



Figure 12. **Success cases** of SVANet on **QuickDraw** dataset. **Green** and **blue** boxes represent **ground truths** and **predictions**, respectively.

B. Additional Qualitative Results

Feature distribution in transfer evaluation. In Fig. 11, we perform transfer evaluation and visualize the feature distribution of SVANet and that of Sketch-DETR using t-SNE [75]. Here, (a) and (b) depict the results in dataset-level transfer, whereas (c) represents the result of category-level transfer. In other words, for (a) and (b), models trained on the **Sketchy** dataset (TU-Berlin dataset for (b)) were employed to map the feature distributions of samples from the **TU-Berlin** dataset (**QuickDraw** dataset for (b)). On the other hand, for (c), models trained on certain categories were utilized to visualize the feature distributions of samples from previously unseen categories. Compared to Sketch-DETR, SVANet appears to be nicely clustered when transferred to

unseen datasets, which implies that SVANet effectively captures class-discriminative representations. When transferred to unseen categories, SVANet embeds the same category into a similar subspace, demonstrating that learned sketch-video mapping can generalize well. Moreover, SVANet shows denser distribution than Sketch-DETR, *i.e.*, only a few data points reach the threshold confidence 0.9 in Sketch-DETR, indicating that our SVANet produces more reliable predictions.

SVOL results: success cases. Fig. 12 shows success cases of SVANet. Our system successfully recognizes the objects that correspond to the query sketch and accurately localizes their bounding boxes in a variety of challenging conditions: (a) various objects appears in a video; (b) multiple object



Figure 13. **Failure cases** of SVANet on QuickDraw dataset. Green and blue boxes represent ground truths and predictions, respectively.

instances with different pose and shape appear in a video; (c) only sketch of a part (face) of object is given; (d) the target objects have different colors; (e) the target objects are occluded by other objects; (f) bad illumination condition.

SVOL results: failure cases. Fig. 13 shows failure cases of SVANet. SVANet suffers particularly when the target object: (a) appears for a very short time (almost 1 or 2 frames out of 32 frames); (b) is too small, and there are numerous distracting factors; (c) is small and moves quickly; (d) is non-salient (here, the target object is a car, but a motorcycle, is detected); (e) is similar to the background.

C. Discussion

C.1. Why Sketch Query?

Sketch query can be more flexible and efficient than language or image query as it allows for more natural and intuitive user input. With sketch query, users can quickly and easily provide a rough sketch of the object they are looking for, rather than having to use specific keywords or search through a pre-existing database of images. This can make it easier for users to find the specific object they are looking for, especially if the image is not easily describable with keywords or if the image does not exist in a pre-existing database. Moreover, sketch query can transcend the language barrier, and can be less prone to ambiguity and errors as the user is able to provide a visual representation of the desired

object. On the other hand, language query requires additional translation when the user’s language changes (*e.g.*, English → French). Additionally, since sketch queries are basically embodiments of real-world objects, the model inherently learns over the visual similarity between the query sketch and the video objects. Therefore, the model can leverage such inductive bias of appearance matching for unseen categories. Sketch query offers a great degree of freedom and can overcome several limitations that other queries (*e.g.*, language or image) may include. There are several more advantages in using sketch query:

1. As the use of touch screen devices (*e.g.*, smartphones, tablets) has recently increased, acquisition of sketch data has become easier.
2. Sketch is not bound by the user’s age, race, and nation, and even those with language difficulties can communicate their thoughts;
3. Our model is effective even with a low-quality sketch (*e.g.*, QuickDraw), therefore users are not required to draw well;

C.2. Limitations

While we believe that focusing on category-level localization in the SVOL problem can take advantage of the abstract nature of sketches, it is also important to consider the limitations of this setting. In this setting, the system may lose nuanced understanding of sketches that could be useful for precisely identifying objects of the same category with different details. For example, if we want to localize a car of a specific make and model, the system may not be able to do so accurately as it is not explicitly taught to differentiate objects within the same category during training. In order to improve the versatility of the SVOL system, future research may investigate on incorporating fine-grained data sources to differentiate objects within the same category.

We also recognize that transfer performance for unseen categories is still far from enough, yet this shows that SVOL is a challenging problem and suggests that better solutions should be found. We hope our findings and analysis will encourage further research in this direction.

C.3. Future Work

We hope future work will develop approaches for the following.

SVOL in large-scale video collection. On an online video platform, users often need to quickly and efficiently find the location of a specific object of interest amid large-scale video collections. In order to be practical in such situations, the SVOL system should be able to retrieve relevant videos, and accurately localize the target objects within the set of retrieved videos. This is similar to the setting for video corpus

moment retrieval [92]. Such a system could greatly enhance the user experience by allowing them to quickly locate the desired object within the video corpus, making the process of finding relevant information faster and more efficient. Although it is beyond the scope of our current work, we believe it to be promising area for future research.

Domain adaptative methods. The significant difference in the appearance and structure between sketches and natural videos poses a challenge for the SVOL system to accurately match them. To alleviate this issue, various domain adaptation techniques [25, 26, 56] can be employed. These techniques aim to align the feature representations of the sketches and natural videos, thus reducing the domain gap. By utilizing these techniques, we anticipate further improvements in the performance of the SVOL system.

Fine-grained SVOL. In this work, we define the SVOL task to be agnostic to shape and pose within the same class, allowing us to localize objects in a video by sketching only key features that are unique and distinctive to that object, such as the ears, eyes, and tails of a cat. This setting has the advantage of being able to identify and locate the object in the video, even if the object’s shape and pose change over time. However, this setting also has its limitations, as it may miss important details that could be useful for differentiating objects within the same class. Therefore, it may be worth exploring a more fine-grained approach to the SVOL problem, by focusing on detailed instance-level information such as shape, pattern, and pose, which can be used to distinguish objects of the same category [91]. This shape- and pose-specific approach, however, also comes with its own challenges. For instance, it may be difficult to match a still sketch to a moving object in a video, as the shape or pose of objects continues to change over time. Thus, to make this approach work effectively, it is essential to have a suitable data pairing that takes into account the dynamic nature of the video. Additionally, further research could explore ways to effectively balance between leveraging the abstract nature of sketches and preserving enough fine-grained details for precise object identification.

On-the-fly SVOL. As opposed to our SVOL setting, which requires a *complete* sketch to be drawn before localization can begin, the “on-the-fly” setting allows for localization to start as soon as the user begins drawing [4]. This approach utilizes each stroke that is drawn in real-time to match it to objects in the video. This allows for sketch-object matching with an *incomplete* sketch (*i.e.*, just a few strokes), which can greatly reduce the time and effort required to draw an accurate sketch. Furthermore, the system can provide immediate feedback based on the ongoing localization results as the user continues to draw, allowing for a more efficient and user-friendly experience. It can help the user to understand how well their sketch is matching with the objects in the video and make adjustments accordingly. This can also make the

task of drawing accurate sketches more manageable for users with less experience or skill.

C.4. Broader Impacts

Our SVANet makes predictions based on learned statistics of the collected dataset, which may reflect biases present in the data, including ones with negative societal impacts. The predictions may not be accurate, thus users exercise caution and should not rely solely on them in real-world applications and it is recommended to use it in conjunction with other forms of analysis and decision-making. Further consideration is warranted regarding this issue.